

Technology 2003

1 Paper Abstract

Name: Dr. Padhraic Smyth

Position/Title: Member of Technical Staff

Affiliation: Jet Propulsion Laboratory, California Institute of Technology.

Address: JPL 238-420, 4800 Oak Grove Drive,
Pasadena, CA 91109.

Phone number: (818) 3543768

FAX number: ((818) 3546825

Govt. Agency/Lab Technology was developed for: NASA (BG-310-20-65-91-00),
ARPA (N00014-92-J-1860)

Category: Test and Measurement

Paper Title: Anomaly Detection using Hidden Markov Models

Description:

Continuous monitoring of complex dynamic systems for the purposes of anomaly or fault detection is an increasingly important issue in diverse areas such as nuclear and chemical plants, manufacturing processes, communications networks and biomedical health monitoring. We have recently developed a novel approach to problems of this nature by combining pattern recognition methods with finite-state hidden Markov models. This approach has general applicability to continuous online monitoring of dynamic systems, with particular potential for applications to onboard monitoring of sensor signals in ground and air vehicles since it is extremely reliable and is computationally very simple to implement. Several recent papers are available which describe both theoretical and practical implementation details [1, 2, 3, 4].

Standard methods for fault detection in dynamic systems have largely relied in the past on two main approaches: control-theoretic methods and heuristic, knowledge-based methods. The control-theoretic approach presumes the existence of an accurate and reliable mathematical system model; frequently in practical applications no such model exists. Knowledge-based methods presume the existence of a domain expert, which is often not the case; furthermore such methods are typically not suited for dealing with time-varying systems with feedback.

An alternative approach which has been proposed in recent years is based on pattern recognition ideas; a classification model is trained on system data to discriminate between normal and abnormal conditions. Since the model adapts to the data from the system, very little prior knowledge is required, making the method very attractive for many practical situations. However, the characteristics of the system which are used as inputs to the classifier (known as the "features" in pattern recognition jargon) can typically be quite noisy (such as features derived from a typical motor current signal in an electro-mechanical system). This can translate into fluctuations in the decisions of the classifier, resulting in an unacceptably large false alarm rate over time.

In recent work at JPL we have developed a novel system which "smooths" such classification estimates over time using the formalism of hidden Markov models. Hidden Markov models have been used with significant success in speech recognition applications for some time; however it is only with our recent work that their general applicability to continuous monitoring has been recognised. Briefly, the system is assumed to follow a Markov trajectory in terms of transiting between normal and abnormal states. The *hidden* aspect of the problem arises due to the fact that these states are not directly observable: only the symptoms (such as sensor time series data) can be directly observed. Hence, the problem is to infer the most likely sequence of system states *given* the symptom data.

In particular, the sampled sensor data is divided into windows over time and from each window a vector of signal parameters is extracted. This data preprocessing step introduces some invariance and robustness with respect to noise. For a given time window, the estimated signal parameter vector provides the input for a standard classification model (we use both kernel density estimators and feedforward neural networks) which in turn produces an instantaneous estimate of the probability of the system's state (conditioned on the particular input). The role of the Markov model is to integrate these probability estimates over time by combining the present state estimate with past state estimates. The hidden Markov model framework provides an elegant theoretical basis to achieve this: essentially it allows the recovery of the most likely sequence of states given the symptom data up to time t . In practice, its most obvious effect is to dramatically reduce the false alarm rate by several orders of magnitude compared to not using any integration over time. The parameters of the hidden Markov model can be shown to be directly related to the overall mean time between failure of the system and hence can be estimated from prior knowledge of system failure rates. The classification part of the model can be trained using system data which has been obtained in advance for this purpose.

We have implemented such a system for online fault detection in JPL's Deep Space Network large ground antennas (34m and 70m dishes) using sampled data signals such as motor currents, position encoders, and tachometer readings. Signal parameters such as standard deviations and fitted autoregressive coefficients are estimated every 4 seconds (an effective downsampling by a factor of 200 from the original sensor sampling rate of 5011 Hz). In various field trials of the system at the Goldstone DSN complex in California, the model has yet to produce a single false alarm and has detected all faults within 8 seconds of occurrence: by comparison, the more standard non-Markov method makes incorrect decisions more than 16% of the time on average (hardware faults are switched into the system in a controlled manner during these tests). The Markov aspect of the model makes the system operationally practical. The success of the method has resulted in a JPL engineering decision that all new DSN antennas will be designed with such a Markov monitoring component built-in. The system will result in a significant reduction in the time spent by DSN operators troubleshooting antenna problems, due to its ability to quickly and accurately detect system abnormalities.

References

1. P. Smyth, 'Hidden Markov models for fault detection in dynamic systems,' submitted to *Pattern Recognition*, February 1993.
2. P. Smyth and J. Mellstrom, 'Deep Space Network antenna monitoring using adaptive time series and hidden Markov models,' accepted for presentation at *AIAA Computing in Aerospace 9 Conference*, San Diego, October 1993 (invited paper).
3. P. Smyth, 'Fault detection in continuous monitoring applications using hidden Markov models,' in *DX-92: Proceedings of the Third International Workshop on the Principles of Diagnosis*, Orcas Island, pp.245-250, October 1992, (invited paper).
4. P. Smyth and J. Mellstrom, 'Detecting novel classes with applications to fault diagnosis,' in *Proceedings of the Ninth International Conference on Machine Learning*, Morgan Kaufmann Publishers: Los Altos, CA, 1992, pp.416-425.